arXiv:1409.7826v1 [math.PR] 27 Sep 2014

GRAPH-BASED PÓLYA'S URN: COMPLETION OF THE LINEAR CASE

YURI LIMA

Abstract. Given a finite connected graph G, place a bin at each vertex. Two bins are called a pair if they share an edge of G. At discrete times, a ball is added to each pair of bins. In a pair of bins, one of the bins gets the ball with probability proportional to its current number of balls. This model was introduced in [BBCL13]. When G is not balanced bipartite, the proportion of balls in the bins converges to a point $w(G)$ almost surely [BBCL13, CL13].

We prove almost sure convergence for balanced bipartite graphs: the possible limit is either a single point $w(G)$ or a closed interval $\mathscr{J}(G)$.

1. Introduction

Let $G = (V, E)$ be a finite connected graph with $V = [m] = \{1, \dots, m\}$ and $|E| = N$, and place a bin at vertex i with $B_i(0) \geq 1$ balls. Consider a random process of adding N balls to the bins at each step, according to the following law: if the numbers of balls after step $n-1$ are $B_1(n-1), \dots, B_m(n-1)$, step n consists of adding, to each edge $\{i, j\} \in E$, one ball either to i or to j, and the probability that the ball is added to i is

$$\mathbb{P}\left[i \text{ is chosen among } \{i, j\} \text{ at step } n\right] = \frac{B_i(n-1)}{B_i(n-1) + B_j(n-1)}\cdot \tag{1.1}$$

We call this model a *graph-based Pólya's urn*.

Call G *bipartite* if there is a partition $V = A \cup B$ such that for every $\{i, j\} \in E$ either $i \in A, j \in B$ or $i \in B, j \in A$. If $\#A = \#B$ we call G *balanced bipartite*, and if $\#A \neq \#B$ we call it *unbalanced bipartite*.

Let $N_0 = \sum_{i=1}^{m} B_i(0)$ be the initial number of balls, let $x_i(n) = \frac{B_i(n)}{N_0 + nN}$, $i \in [m]$, and let $x(n) = (x_1(n), \dots, x_m(n))$, the proportion of balls in the bins after step n.

Theorem 1.1. *If G is a finite, connected, balanced bipartite graph, then there is a closed interval $\mathscr{J} = \mathscr{J}(G)$ such that $x(n)$ converges to a point of $\mathscr{J}$ almost surely.*

In some cases $\mathscr{J}$ is a singleton. When it is not, the point to which $x(n)$ converges depends on the realization of the process $x(n)$.

Graph-based Pólya's urns were introduced in [BBCL13]: for a fixed $\alpha > 0$, one of the bins gets the ball with probability proportional to the α power of its current number of balls. When $\alpha = 1$ the model is (1.1), hence we call it the *linear case*.

Date: September 30, 2014.

2010 *Mathematics Subject Classification.* Primary: 60K35. Secondary: 37C10.

Key words and phrases. Gradient-like system, Pólya's urn, reinforcement, stochastic approximation algorithms.

The author was supported by the Brin Fellowship.

Graph-based Pólya's urns extend the classical Pólya's urn and many of its variants, see [Pem07]. E.g. if G is the complete graph with m vertices then the model is a Pólya's urn with m colors.

$\mathscr{J}$ depends on the structure of the graph. If G is not bipartite then $\mathscr{J}$ is a singleton [BBCL13]. This was extended to unbalanced bipartite graphs [CL13]. The remaining case, when G is balanced bipartite, was conjectured in [CL13, Conjecture 5.4]. Theorem 1.1 confirms it, and completes the description of possible limits.

Corollary 1.2. *If G is a finite, connected graph, then there is a closed interval $\mathscr{J} = \mathscr{J}(G)$ such that $x(n)$ converges to a point of $\mathscr{J}$ almost surely. If G is not balanced bipartite, then $\mathscr{J}$ is a singleton.*

Additionally to being natural generalizations of Pólya's urns, graph-based Pólya's urns model some competing networks [BBCL13]: Imagine there are 3 companies, denoted by M, A, G. Each company sells two products. M sells OS and SE, A sells OS and SP, G sells SE and SP. Each pair of companies compete on one product. The companies try to use their global size and reputation to boost sales. Which company will sell more products in the long term? The interaction between the companies form a triangular network: a vertex represents a company and an edge represents a product. Under further simplifications, graph-based Pólya's urns describe in broad strokes the long-term evolution of such competition.

Another example comes from a repeated game in which agents improve their skill by gaining experience [SP00]. The interaction network between agents is modeled by a graph. At each round a pair is competing for a ball. A competitor improves his skill with time, and the number of balls in his bin represents his skill level. See [BBCL13, §1] and references therein for more applications.

The sequence $x(n)$ is a *stochastic approximation algorithm*. These are small perturbations of a vector field. In many cases, there is a relation between the limit set of $x(n)$ and the equilibria of the vector field. For graph-based Pólya's urns, the vector field is gradient-like [BBCL13, Lemma 4.1], thus the limit set of $x(n)$ is almost surely contained in the equilibria set of the vector field.

Since limit sets are connected, if there are finitely many equilibria then $x(n)$ converges almost surely to some equilibrium. Some of them are unstable, and some are not (see §2). The probability that $x(n)$ converges to an unstable equilibrium is zero [BBCL13, Lemma 5.2], see also Theorem 3.1 here. Hence at least one equilibrium is non-unstable. Complementary to this, non-unstable equilibria generate Lyapunov functions [CL13, Lemmas 3.1 and 3.2]. If G is not balanced bipartite, this implies that there is at most one non-unstable equilibrium. Combined, these two arguments imply the second part of Corollary 1.2 [CL13, Theorem 1.1].

If there are infinitely many non-unstable equilibria, then $x(n)$ could wander around without converging to any of them. To prove convergence, one needs to understand the attracting/repelling properties of the equilibria. E.g. if G is regular and balanced bipartite then the set of non-unstable equilibria is an interval and the eigenvalues in transverse directions have negative real part [BBCL13, Lemma 10.1], thus $x(n)$ converges almost surely to a point of the interval [CL13, Theorem 1.2]. Here is a heuristic explanation: the orbits of the vector field converge exponentially fast to the interval, and the random model converges exponentially fast to its limit set [CL13, Lemma 4.1]. This prevents $x(n)$ of wandering around the interval.

We prove that for balanced bipartite graphs the set of non-unstable equilibria is an interval $\mathscr{J}$, possibly reduced to a point. When it is not a point, we prove that all eigenvalues are real, and those in transverse directions to $\mathscr{J}$ are negative. Under these conditions, we apply the methods of [CL13, Theorem 1.2] to prove that $x(n)$ converges to a point of $\mathscr{J}$ almost surely. This gives Theorem 1.1.

2. Stochastic approximation algorithms

Graph-based Pólya's urn are an example of stochastic approximation algorithms [BBCL13]. In this section we recall some previous results of [BBCL13, CL13] and explain how a graph-based Pólya's urn is related to a vector field.

Stochastic approximation algorithm: A *stochastic approximation algorithm* is a discrete time process $\{x(n)\}_{n\geq 0} \subset \mathbb{R}^m$ of the form

$$x(n+1) - x(n) = \gamma_n \left[F(x(n)) + u_n\right] \tag{2.1}$$

where $\{\gamma_n\}_{n\geq 0}$ is a sequence of nonnegative scalar gains, $F : \mathbb{R}^m \to \mathbb{R}^m$ is a vector field, and $u_n \in \mathbb{R}^m$ is a random vector that depends on $x(n)$ only.

Let $\mathscr{F}_n$ be the sigma-algebra generated by the process up to step n. Since u_n only depends on $x(n)$ we can assume, after changing F, that $\mathbb{E}\left[u_n|\mathscr{F}_n\right] = 0$.

Graph-based Pólya's urns are stochastic approximation algorithms with $\gamma_n = \frac{1}{\frac{N_0}{N}+(n+1)}$ and vector field F defined by the equations:

$$\left\{\begin{array}{rcl} \dfrac{dv_1}{dt} &=& -v_1 + \dfrac{1}{N}\displaystyle\sum_{j\sim 1} \frac{v_1}{v_1+v_j} \\ &\vdots& \\ \dfrac{dv_m}{dt} &=& -v_m + \dfrac{1}{N}\displaystyle\sum_{j\sim m} \frac{v_m}{v_m+v_j}\,. \end{array}\right. \tag{2.2}$$

See [BBCL13, §3.2].

Domain of F: Fix $c < \frac{1}{N}$, and let Δ be the set of vectors $(v_1, \ldots, v_m) \in \mathbb{R}^m_{\geq 0}$ with $\sum_{i=1}^m v_i = 1$ and $v_i + v_j \geq c$ for all $\{i,j\} \in E$. The vector field $F : \Delta \to T\Delta$ is Lipschitz, and it induces a semiflow [BBCL13, Lemma 2.1].

The vector field F is gradient-like. This was proved in [BBCL13, Lemma 4.1].

Equilibria set: $v \in \Delta$ is called an *equilibrium* if $F(v) = 0$. v is called *unstable* if $DF(v)$ has an eigenvalue with negative real part, and *non-unstable* otherwise. The *equilibria set* is $\Lambda = \{v \in \Delta : v \text{ is equilibrium}\}$.

Lyapunov function: Let $U \subseteq \Delta$. A continuous map $L : \Delta \to \mathbb{R}$ is called a *Lyapunov function for U* if it is strictly monotone along any integral curve of F outside U. If $U = \Lambda$, we call L a *strict Lyapunov function* and F *gradient-like*.

Let $L : \Delta \to \mathbb{R}$ be the function

$$L(v_1, \ldots, v_m) = -\sum_{i=1}^m v_i + \frac{1}{N}\sum_{\{i,j\}\in E} \log\left(v_i + v_j\right). \tag{2.3}$$

L is a strict Lyapunov function for F: because $\frac{dv_i}{dt} = v_i \frac{\partial L}{\partial v_i}$, then

$$\frac{d}{dt}(L \circ v) = \sum_{i=1}^{m} \frac{\partial L}{\partial v_i} \frac{dv_i}{dt} = \sum_{i=1}^{m} v_i \left(\frac{\partial L}{\partial v_i} \right)^2 \geq 0.$$

Equality holds iff $v_i \frac{\partial L}{\partial v_i} = 0$ for all i iff $v \in \Lambda$.

We divide the singularities according to the faces of Δ. Given $S \subseteq [m]$, let $\Delta_S = \{v \in \Delta : v_i = 0 \text{ iff } i \notin S\}$. The restriction $F \restriction_{\Delta_S}$ is a semiflow. Let $\Lambda_S = \{v \in \Delta_S : \frac{\partial L}{\partial v_i}(v) = 0, \forall i \in S\}$. A direct calculation shows that $\Lambda = \bigcup_{S \subseteq [m]} \Lambda_S$ [BBCL13, Lemma 2.1]. Because L is a concave function, so is $L \restriction_{\Delta_S}$, hence Λ_S is the set of maxima of $L \restriction_{\Delta_S}$.

Relation between $\{x(n)\}_{n\geq 0}$ and F. Let $\{\Phi_t\}_{t\geq 0}$ be the semiflow induced by F. Let $\tau_n = \sum_{i=0}^{n} \gamma_i$, and let $\{X(t)\}_{t\geq 0}$ be the interpolation of $\{x(n)\}_{n\geq 0}$: $X(\tau_n) = x(n)$ and $X \restriction_{[\tau_n, \tau_{n+1}]}$ is linear. Let d be the euclidean distance on Δ.

Theorem 2.1. [BBCL13, CL13] *The limit set of $x(n)$ is contained in Λ almost surely, and*

$$\sup_{T>0} \limsup_{t\to+\infty} \frac{1}{t} \log \left(\sup_{0\leq h\leq T} d\big(X(t+h), \Phi_h(X(t))\big) \right) \leq -\frac{1}{2}. \tag{2.4}$$

The first part was proved in [BBCL13, §3.1 and §3.2]. It is an application of the general theory of stochastic approximation algorithms [Ben96, Ben99]. The second part is [CL13, Lemma 4.1]. It follows from shadowing techniques that relate the speed of convergence of the interpolated process and the vector field [Ben99, Prop. 8.3]. The right hand side of the inequality is the log-convergence rate $\frac{1}{2} \limsup \frac{\log \gamma_n}{\tau_n}$, which for graph-based Pólya's urns equals $-\frac{1}{2}$.

3. Unstable and non-unstable equilibria

Write $F = (F_1, \ldots, F_m)$, $F_i = v_i \frac{\partial L}{\partial v_i}$. Fix $w \in \Lambda_S$, and let $DF(w) : T_w\Delta \to T_w\Delta$. In coordinates $v_1, \ldots, v_m$, $DF(w)$ equals the jacobian matrix $JF(w) = \left(\frac{\partial F_i}{\partial v_j} \right)_{i,j}$:

$$\frac{\partial F_i}{\partial v_j} = \begin{cases} v_i \dfrac{\partial^2 L}{\partial v_i \partial v_j} & \text{if } i \neq j, \\ \dfrac{\partial L}{\partial v_i} + v_i \dfrac{\partial^2 L}{\partial v_i^2} & \text{if } i = j. \end{cases} \tag{3.1}$$

Without loss of generality, assume that $S = \{k+1, \ldots, m\}$. Thus

$$JF(w) = \begin{bmatrix} A & 0 \\ C & B \end{bmatrix} \tag{3.2}$$

where A is a $k \times k$ diagonal matrix with $a_{ii} = \frac{\partial L}{\partial v_i}(w)$, $i \in [k]$.

Non-convergence to unstable equilibria. The spectrum of $JF(w)$ is the union of the spectra of A and B. Introduce the inner product $(x,y) = \sum_{i=k+1}^{m} x_i y_i / v_i$. B is self-adjoint and negative semidefinite (by the concavity of L), hence the eigenvalues of B are real and nonpositive. Therefore $JF(w)$ has a real positive eigenvalue iff $a_{ii} > 0$ for some $i \in [k]$. In summary:

$$w \in \Lambda_S \text{ is unstable } \iff \exists i \notin S \text{ s.t. } \tfrac{\partial L}{\partial v_i}(w) > 0. \tag{3.3}$$

Theorem 3.1. [BBCL13, Lemma 5.2] *If w is an unstable equilibrium, then*

$$\mathbb{P}\left[\lim_{n\to\infty} x(n) = w\right] = 0.$$

In particular, if Λ is finite then $x(n)$ almost surely converges to a non-unstable equilibrium. The proof is probabilistic and follows the lines of [Pem92, §3 and §4], see also [Ben99, §9].

Non-unstable equilibria and Lyapunov functions. Let $w \in \Lambda_S$ non-unstable. By (3.3), $\frac{\partial L}{\partial v_i}(w) \le 0$ for every $i \notin S$. Since $\frac{\partial L}{\partial v_i}(w) = 0$ for $i \in S$, we have [CL13]:

$$w \in \Lambda_S \text{ is non-unstable } \iff \tfrac{\partial L}{\partial v_i}(w) = 0 \ \ \forall i \in S, \text{ and } \tfrac{\partial L}{\partial v_i}(w) \le 0 \ \ \forall i \notin S. \tag{3.4}$$

In particular, every $w \in \Lambda_{[m]}$ is non-unstable.

For every non-unstable equilibrium there is a Lyapunov function that gives extra information on the convergence of the vector field [CL13]. This fact will be used in §4 and §5, thus we state it in a general form. Given $w \in \Delta_S$ and $\chi \in (0, \min_{i\in S} w_i]$, let $\Delta^{w,\chi} = \{v \in \Delta : v_i \ge \chi, \forall i \in S\}$ (we do not require that $v_i = 0$ for $i \notin S$). $\Delta^{w,\chi}$ is a closed convex set that contains w at its boundary. The next result is a summary of [CL13, Lemmas 3.1 and 3.2].

Lemma 3.2. *Let $w \in \Lambda_S$ non-unstable. Then there is a closed interval $J = J(w,\chi)$ such that $H : \Delta^{w,\chi} \to \mathbb{R}$, $H(v) = \sum_{i\in S} w_i \log v_i$, is a Lyapunov function for J.*

In particular, every orbit of $F \restriction_{\Delta^{w,\chi}}$ converges to J.

Proof. Inside $\Delta^{w,\chi}$ the function H is differentiable, and

$$\frac{d}{dt}(H \circ v) = \sum_{i\in S} w_i \frac{1}{v_i}\frac{dv_i}{dt} = \sum_{i\in S} w_i \frac{\partial L}{\partial v_i} = \sum_{i=1}^{m} w_i \frac{\partial L}{\partial v_i} = -1 + \frac{1}{N}\sum_{\{i,j\}\in E} \frac{w_i + w_j}{v_i + v_j}.$$

Let $f : \Delta^{w,\chi} \to \mathbb{R}$, $f(v) = -1 + \frac{1}{N}\sum_{\{i,j\}\in E} \frac{w_i+w_j}{v_i+v_j}$. Observe that $f(w) = 0$. We will show that $f(v) \ge 0$, with equality iff $v \in J$ (to be defined below).

Step 1: f is convex.

Since $x > 0 \mapsto \frac{1}{x}$ is convex, each $v \in \Delta^{w,\chi} \mapsto \frac{w_i+w_j}{v_i+v_j}$ is convex. Thus f is the sum of convex functions.

Step 2: w is a global minimum of f.

Since f is convex, it is enough to prove that w is a local minimum of f. Let $v = w + (\varepsilon_1, \dots, \varepsilon_m)$ with $\varepsilon_1, \dots, \varepsilon_m$ small enough. Of course, $\varepsilon_i \ge 0$ for $i \notin S$.

Applying the inequality $\frac{x}{x+\varepsilon} - 1 \geq -\frac{\varepsilon}{x}$ for $x, x+\varepsilon > 0$, we have

$$\begin{aligned} f(v) - f(w) &= \frac{1}{N} \sum_{\{i,j\} \in E} \left[\frac{w_i + w_j}{v_i + v_j} - 1 \right] \\ &\geq \frac{1}{N} \sum_{\{i,j\} \in E} -\frac{\varepsilon_i + \varepsilon_j}{w_i + w_j} \\ &= -\sum_{i=1}^{m} \varepsilon_i \frac{1}{N} \sum_{j \sim i} \frac{1}{w_i + w_j} \\ &= -\sum_{i=1}^{m} \varepsilon_i \left[1 + \frac{\partial L}{\partial v_i}(w) \right] \\ &= -\sum_{i=1}^{m} \varepsilon_i \frac{\partial L}{\partial v_i}(w) \\ &\geq 0, \end{aligned}$$

since $\varepsilon_i \frac{\partial L}{\partial v_i}(w) = 0$ for $i \in S$, and $\varepsilon_i \frac{\partial L}{\partial v_i}(w) \leq 0$ for $i \notin S$. Hence w is a local minimum of f.

Step 3: The set of global minima of f is a closed interval $J \ni w$.

The set of global minima of a convex function is convex. Thus if $v \in \Delta^{w,\chi}$ with $f(v) = f(w)$ then $f(tv + (1-t)w) = tf(v) + (1-t)f(w)$ for all $t \in [0,1]$. Because $x > 0 \mapsto \frac{1}{x}$ is strictly convex, we get $v_i + v_j = w_i + w_j$ for all $\{i,j\} \in E$, i.e.

$$v_i - w_i = -(v_j - w_j), \ \forall \{i,j\} \in E. \tag{3.5}$$

We divide the analysis of (3.5) into three cases:

- G is not bipartite: G has an odd cycle, thus (3.5) implies $v = w$. Take $J = \{w\}$.
- G is unbalanced bipartite: let $V = A \cup B$ be the bipartition, $\#A \neq \#B$. By (3.5), there is $\eta \in \mathbb{R}$ such that

$$v_i = \begin{cases} w_i + \eta & \text{if } i \in A, \\ w_i - \eta & \text{if } i \in B. \end{cases} \tag{3.6}$$

  Summing up on i, we get $\eta(\#A - \#B) = 0 \Rightarrow \eta = 0$. Take $J = \{w\}$.
- G is balanced bipartite: let $V = A \cup B$ be the bipartition, $\#A = \#B$. As in the previous case, (3.6) holds. Take $J = \{v \in \Delta^{w,\chi} : v \text{ satisfies (3.6)}\}$. J is a closed interval, and $f \restriction_J$ is identically zero.

□

We want to avoid the dependence of J on w, χ.

The interval $\mathscr{J}$: $\mathscr{J}$ is the maximal extension of J to Δ.

$\mathscr{J}$ is an interval whose endpoints belong to $\partial\Delta$, one of which is w, and whose interior is contained in $\Delta_{[m]}$. Furthermore:

(i) $\mathscr{J}$ is uniquely determined by any of its points.
(ii) $\frac{\partial L}{\partial v_i} \restriction_{\mathscr{J}}$ is constant and equal to $\frac{\partial L}{\partial v_i}(w)$ for all i, because of (3.6).

4. Not balanced bipartite graphs

If G is not balanced bipartite, then Corollary 1.2 holds with $\mathscr{J}$ = singleton [CL13, Theorem 1.1]. We include the proof for completeness.

Step 1: L is strictly concave.

We have $L(tv + (1-t)w) \geq tL(v) + (1-t)L(w)$ for all $v, w \in \Delta, t \in [0,1]$. Equality holds iff (3.5) holds iff $v = w$, because:

- If G is not bipartite then it has an odd cycle, hence $v = w$.
- If G is unbalanced bipartite then (3.6) holds, hence $v = w$.

Step 2: Λ is finite.

$L \restriction_{\Delta_S}$ is strictly concave, because it is the restriction of L to a convex set. Thus Λ_S is either empty or a singleton, and $\Lambda = \bigcup_{S \subseteq [m]} \Lambda_S$ is finite.

Step 3: There is at least one non-unstable equilibrium.

This follows directly from Theorem 3.1.

Step 4: There is at most one non-unstable equilibrium.

Suppose $w \neq \widetilde{w}$ are non-unstable equilibria. Let $H : \Delta^{w,\chi} \to \mathbb{R}$, $\widetilde{H} : \Delta^{\widetilde{w},\chi} \to \mathbb{R}$ as in Lemma 3.2. Take $\chi > 0$ small enough such that $\Delta^{w,\chi} \cap \Delta^{\widetilde{w},\chi} \neq \emptyset$. Every orbit of F starting from $\Delta^{w,\chi} \cap \Delta^{\widetilde{w},\chi}$ converges simultaneously to w and $\widetilde{w}$, a contradiction.

By steps 3 and 4 there is a unique non-unstable equilibrium $w = w(G)$, and $x(n)$ converges to w almost surely.

5. Balanced bipartite graphs

Let $V = A \cup B$ be the bipartition, $\#A = \#B$. We consider two cases.

First case: $\Lambda_{[m]} = \emptyset$. Steps 1–3 below are in [CL13, Corollary 5.2].

Step 1: $L \restriction_{\Delta_S}$ is strictly concave for every $S \neq [m]$.

If $L(tv + (1-t)w) = tL(v) + (1-t)L(w)$ with $v, w \in \Delta_S, t \in [0,1]$, then (3.6) holds. For $i \in [m] \backslash S$ we have $v_i = w_i = 0$, hence $\eta = 0$.

Step 2: Λ is finite.

By step 1, if $S \neq [m]$ then Λ_S is either empty or a singleton. Since $\Lambda_{[m]} = \emptyset$, $\Lambda = \bigcup_{S \subseteq [m]} \Lambda_S$ is finite.

Step 3: There is at least one non-unstable equilibrium.

Again, this is consequence of Theorem 3.1.

Step 4: There is at most one non-unstable equilibrium.

Let $w \neq \widetilde{w}$ be non-unstable equilibria, let $\Delta^{w,\chi}, \Delta^{\widetilde{w},\chi}$ as in Lemma 3.2, and $\mathscr{J}, \widetilde{\mathscr{J}}$ be the maximal intervals defined at the end of §3. Choose $\chi > 0$ small enough so that $\Delta^{w,\chi} \cap \Delta^{\widetilde{w},\chi} \neq \emptyset$. Every orbit of F starting from $\Delta^{w,\chi} \cap \Delta^{\widetilde{w},\chi}$ converges to both $\mathscr{J}$ and $\widetilde{\mathscr{J}}$. Since F is gradient-like they also converge to Λ, thus $\mathscr{J} \cap \widetilde{\mathscr{J}} \cap \Lambda \neq \emptyset$. This will give the contradiction we are looking for.

Since $\mathscr{J} \cap \widetilde{\mathscr{J}} \neq \emptyset$ and $\mathscr{J}, \widetilde{\mathscr{J}}$ are determined by any of its points, $\mathscr{J} = \widetilde{\mathscr{J}}$. w is an endpoint of $\mathscr{J}$, and $\widetilde{w}$ is and endpoint of $\widetilde{\mathscr{J}}$, thus $w, \widetilde{w}$ are the two

endpoints of $\mathscr{J} = \widetilde{\mathscr{J}}$. In particular, if $w_i = 0$ then $\widetilde{w}_i > 0$. This gives that $([m]\backslash S) \cap ([m]\backslash \widetilde{S}) = \emptyset$, hence $S \cup \widetilde{S} = V$. By (3.4) we get $\mathscr{J} \subseteq \Lambda$: if $v \in \mathscr{J}$ then $\frac{\partial L}{\partial v_i}(v) = \frac{\partial L}{\partial v_i}(w) = 0$ for $i \in S$, and $\frac{\partial L}{\partial v_i}(v) = \frac{\partial L}{\partial v_i}(\widetilde{w}) = 0$ for $i \in \widetilde{S}$. In particular $\emptyset \neq \text{int}(\mathscr{J}) \subset \Lambda_{[m]}$, a contradiction.

By steps 3 and 4, there is a unique non-unstable equilibrium $w = w(G)$, and $x(n)$ converges to w almost surely.

Second case: $\Lambda_{[m]} \neq \emptyset$. We will prove that there is a non-degenerate interval $\mathscr{J} = \mathscr{J}(G)$ such that $x(n)$ converges to a point of $\mathscr{J}$ almost surely.

Step 1: The set on non-unstable equilibria is a closed interval $\mathscr{J}$.

Remember that any $w \in \Lambda_{[m]}$ is non-unstable, since $\frac{\partial L}{\partial v_i}(w) = 0$ for all i. Apply Lemma 3.2 to w, and let $\mathscr{J}$ be the maximal interval defined as in the end of §3. $\frac{\partial L}{\partial v_i} \restriction_{\mathscr{J}}$ is identically zero for all i, hence $\mathscr{J}$ is an interval of non-unstable equilibria.

We now show that $\mathscr{J}$ is *the* set of all non-unstable equilibria. The proof is similar to the proof of step 4 of the first case. Let $\widetilde{w}$ be a non-unstable equilibrium, and let $\widetilde{\mathscr{J}}$ be the maximal interval defined as in the end of §3. If $\chi > 0$ is sufficiently small then $J \cap \widetilde{J} \neq \emptyset$, thus $\mathscr{J} = \widetilde{\mathscr{J}}$. Hence $\widetilde{w} \in \mathscr{J}$.

Remark 5.1. Step 1 above and the first case characterize, for balanced bipartite graphs, when $\mathscr{J}$ is a singleton or not.

- $\mathscr{J}$ is a singleton iff there is a non-unstable equilibrium w with $\frac{\partial L}{\partial v_i} < 0$ for some i: otherwise w would define an interval of equilibria whose interior is a subset of $\Lambda_{[m]}$.
- $\mathscr{J}$ is a non-degenerate interval iff $\frac{\partial L}{\partial v_i} = 0$, $i \in [m]$, for all non-unstable equilibria: since $\frac{\partial L}{\partial v_i} \restriction_{\Lambda_{[m]}} \equiv 0$ and $\frac{\partial L}{\partial v_i} \restriction_{\mathscr{J}}$ is constant, we have $\frac{\partial L}{\partial v_i} \restriction_{\mathscr{J}} \equiv 0$.

We will make use of this in the discussion of some examples, see §6.

Step 2: If $w \in \text{int}(\mathscr{J})$ then all eigenvalues of $DF(w)$ are real, and any eigenvalue in a transverse direction to $\mathscr{J}$ is negative.

This was proved for *regular* balanced bipartite graphs [BBCL13, Lemma 10.1]. The question remained open for a general balanced bipartite graph.

Let $w \in \text{int}(\mathscr{J})$, thus $\frac{\partial L}{\partial v_i}(w) = 0$ for all i. By (3.1), $DF(w)$ is the restriction of the matrix $B = \left(v_i \frac{\partial^2 L}{\partial v_i \partial v_j}\right)$ to $T_w\Delta$. Let $A = \left(\frac{\partial^2 L}{\partial v_i \partial v_j}\right)$ be the Hessian of L in the coordinates $v_1, \ldots, v_m$. The rows of B are positive multiples of the rows of A.

The matrix A is symmetric, thus its eigenvalues are real. Since $\text{int}(\mathscr{J})$ is the set of global maxima of $L \restriction_{\Delta_{[m]}}$, A is negative semidefinite and zero is a simple eigenvalue, i.e. every eigenvalue in a transverse direction to $\mathscr{J}$ is negative. We claim that the same is true for B. Remind the inner product $(x, y) = \sum_{i=1}^m \frac{x_i y_i}{v_i}$ introduced in §3, and let $\langle x, y\rangle = \sum_{i=1}^m x_i y_i$ be the canonical inner product.

Since $(Bx, y) = \langle Ax, y\rangle$, B is self-adjoint: $(Bx, y) = \langle Ax, y\rangle = \langle x, Ay\rangle = (x, By)$. Thus the eigenvalues of B are real. Let us prove that one of them is zero and the others are negative.

- 0 is a simple eigenvalue: $B = DA$, where D is the diagonal matrix with diagonal entries $v_1, \dots, v_m$. Since $v \in \Delta_{[m]}$, D is invertible, thus the kernels of A and B coincide. In particular, the kernel of B is one-dimensional.
- 0 is the largest eigenvalue: let $M = \max_{i\in[m]} v_i > 0$, thus $(x,x) \geq M^{-1}\langle x, x\rangle$. Let $\lambda_1(\cdot)$ denote the largest eigenvalue of a matrix. By the variational characterization of eigenvalues of hermitian matrices (see [HJ13, Theorem 4.2.2]),

$$\lambda_1(B) = \max_{x\neq 0} \frac{(Bx,x)}{(x,x)} \leq M \max_{x\neq 0} \frac{\langle Ax, x\rangle}{\langle x, x\rangle} = M\lambda_1(A) = 0.$$

This concludes the proof of step 2.

STEP 3: $x(n)$ converges to a point of $\mathscr{J}$ almost surely.

It is enough to prove that the interpolated orbits $X(t)$ converge to a point of $\mathscr{J}$ almost surely. This is true for regular balanced bipartite graphs [CL13, Theorem 1.2]. Here is a heuristic of the proof: since the interpolated process converges exponentially fast (Theorem 2.1) and the orbits of F also converge exponentially fast (step 2), the interpolated process cannot wander around $\mathscr{J}$. Provided these are true, the proof in [CL13] applies ipsis literis. We include it for completeness.

For a fixed closed interval $I \subset \mathrm{int}(\mathscr{J})$, and a small neighborhood U of I in Δ, there is a foliation $\{\mathscr{F}_x\}_{x\in U}$ such that:

- $\mathscr{F}_x$ is a submanifold with $\mathscr{F}_x \pitchfork \mathscr{J}$ at a single point $\pi(x)$.
- $\pi(x)$ is a hyperbolic attractor for $F \restriction_{\mathscr{F}_x}$. The speed of convergence depends on the negative eigenvalues of $DF(\pi(x))$.

This is an application of the theory of invariant manifolds for normally hyperbolic sets, see [HPS77, Theorem 4.1].

The map $\pi : U \to \mathscr{J}$ is not necessarily a projection (it is not even linear), but since $\mathscr{F}_x$ depends smoothly on x, if U is small enough then π is 2–Lipschitz:

$$d(\pi(x), \pi(y)) \leq 2d(x,y), \forall\, x, y \in U. \tag{5.1}$$

Fix a small parameter $\varepsilon > 0$ and reduce U, if necessary, so that

$$U = \{x \in \Delta : \pi(x) \in I \text{ and } d(x, \pi(x)) < \varepsilon\}. \tag{5.2}$$

Let $c = \max\{\lambda : \lambda \neq 0 \text{ is eigenvalue of } DF(x), x \in I\}$. By step 2, $c < 0$. Thus there is $K > 0$ such that

$$d(\Phi_t(x), \pi(x)) \leq Ke^{ct} d(x, \pi(x)), \forall\, x \in U, \forall\, t \geq 0. \tag{5.3}$$

(Remind: $\{\Phi_t\}_{t\geq 0}$ is the semiflow induced by F.)

Fix an interpolated orbit $X(t)$ that does not converge to the endpoints of $\mathscr{J}$. It has an accumulation point in $\mathrm{int}(\mathscr{J})$. Let $I \subset \mathrm{int}(\mathscr{J})$ be an interval containing such point, and let U as in (5.2).

Lemma 5.2. [CL13, Lemma 4.4] *Assume that $X(t) \in U$. If t, T are large enough, then*

(i) $d(\pi(X(t+T)), \pi(X(t))) < 2e^{-\frac{t}{4}}$.
(ii) $X(t+T) \in U$.

Proof. To simplify the notation, denote $X(t)$ by X and $X(t+T)$ by $X(T)$.

(i) Since $\pi(\Phi_T(X)) = \pi(X)$ and π is 2–Lipschitz,

$$d(\pi(X(T)), \pi(X)) = d(\pi(X(T)), \pi(\Phi_T(X))) \leq 2d(X(T), \Phi_T(X)).$$

By (2.4), $d(X(T), \Phi_T(X)) < e^{-\frac{t}{4}}$ for large t, therefore $d(\pi(X(T)), \pi(X)) < 2e^{-\frac{t}{4}}$ for large t.

In particular, $\pi(X(T)) \in I$ for large t.

(ii) Since $\pi(X(T)) \in I$, it remains to estimate $d(X(T), \pi(X(T)))$. By the triangular inequality, (5.1) and (5.3), we have

$$\begin{aligned} d(X(T), \pi(X(T))) &\leq d(X(T), \Phi_T(X)) + d(\Phi_T(X), \pi(\Phi_T(X))) \\ &\quad + d(\pi(\Phi_T(X)), \pi(X(T))) \\ &\leq 3d(X(T), \Phi_T(X)) + d(\Phi_T(X), \pi(X)) \\ &\leq 3e^{-\frac{t}{4}} + Ke^{cT} d(X, \pi(X)) \\ &\leq 3e^{-\frac{t}{4}} + Ke^{cT}\varepsilon \\ &< \varepsilon \end{aligned}$$

provided $3e^{-\frac{t}{4}} < \frac{\varepsilon}{2}$ and $Ke^{cT} < \frac{1}{2}$. □

The second part of the lemma allows us to apply it inductively to the points $X_k := X(t + kT), k \geq 0$. For that, choose t, T large enough so that $2\sum_k e^{-\frac{t+kT}{4}} < d(\pi(X_0), \mathscr{I}\backslash I)$. By Lemma 5.2, if $X_k \in U$ then $X_{k+1} \in U$ and $d(\pi(X_{k+1}), \pi(X_k)) < 2e^{-\frac{t+kT}{4}}$. Thus $\pi(X_k)$ converges, say $\lim \pi(X_k) = \widetilde{x}$.

The proof of Lemma 5.2(ii) also gives that

$$d(X_k, \pi(X_k)) \leq 3e^{-\frac{t+(k-1)T}{4}} + Ke^{cT} d(X_{k-1}, \pi(X_{k-1})),\ k \geq 1.$$

Let $\lambda = Ke^{cT}$, thus:

$$\begin{aligned} d(X_k, \pi(X_k)) &\leq 3e^{-\frac{t}{4}} \left(e^{-\frac{(k-1)T}{4}} + \lambda e^{-\frac{(k-2)T}{4}} + \cdots + \lambda^{k-1} \right) + \lambda^k d(X_0, \pi(X_0)) \\ &\leq 3e^{-\frac{t}{4}} k \left(\max\left\{ e^{-\frac{T}{4}}, \lambda \right\} \right)^{k-1} + \lambda^k d(X_0, \pi(X_0)). \end{aligned}$$

When T is large, $\max\{e^{-\frac{T}{4}}, \lambda\} < 1$, hence $\lim d(X_k, \pi(X_k)) = 0$. Since $d(X_k, \widetilde{x}) \leq d(X_k, \pi(X_k)) + d(\pi(X_k), \widetilde{x})$, it follows that $\lim X_k = \widetilde{x} \in I$.

Now let $s \in [t + kT, t + (k+1)T)$. By the triangular inequality and (2.4)

$$\begin{aligned} d(X(s), \widetilde{x}) &= d(X(s), \Phi_{s-(t+kT)}(\widetilde{x})) \\ &\leq d(X(s), \Phi_{s-(t+kT)}(X_k)) + d(\Phi_{s-(t+kT)}(X_k), \Phi_{s-(t+kT)}(\widetilde{x})) \\ &\leq e^{-\frac{t+kT}{4}} + c(T) d(X_k, \widetilde{x}), \end{aligned}$$

where $c(T) > 0$ is the supremum of the Lipschitz constants of $\Phi_\delta, \delta \in [0, T]$. Therefore $X(t)$ converges to $\widetilde{x}$.

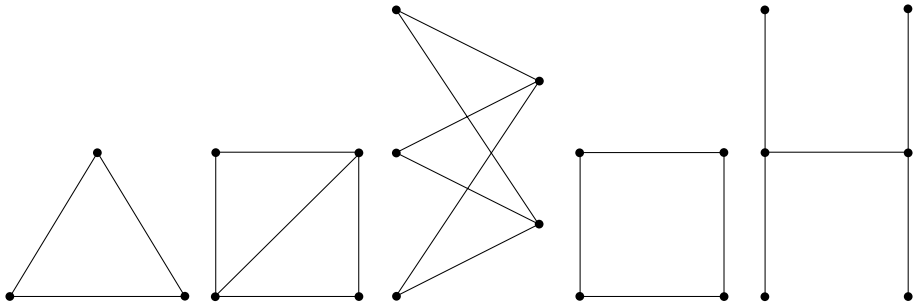

FIGURE 1. (a)–(b) are not bipartite; (c) is unbalanced bipartite; (d)–(e) are balanced bipartite.

6. Concluding remarks

Some examples. Consider the graphs in Figure 1. We show that all cases considered in the proof of Corollary 1.2 occur. Remind: every $v \in \Lambda_{[m]}$ is non-unstable.

(a) The triangle is not bipartite and $(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}) \in \Lambda_{[m]}$, thus $\mathscr{J} = \{(\frac{1}{3}, \frac{1}{3}, \frac{1}{3})\}$. Similarly, complete graphs and cycles of odd length satisfy $\mathscr{J}$ = uniform distribution.

(b) The graph is not bipartite and $(0, \frac{1}{2}, 0, \frac{1}{2})$ is a non-unstable equilibrium (since $\frac{\partial L}{\partial v_1} = -\frac{1}{5}$), thus $\mathscr{J} = \{(0, \frac{1}{2}, 0, \frac{1}{2})\}$.

(c) The graph is unbalanced bipartite and $(0, 0, 0, \frac{1}{2}, \frac{1}{2})$ is a non-unstable equilibrium (since $\frac{\partial L}{\partial v_1} = -\frac{1}{3}$), thus $\mathscr{J} = \{(0, 0, 0, \frac{1}{2}, \frac{1}{2})\}$. More generally, if $K_{i,j}$ is the complete bipartite graph and if $i > j$, then $\mathscr{J} = \{(0, \dots, 0, \frac{1}{j}, \dots, \frac{1}{j})\}$.

(d) The square is balanced bipartite and $(\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4}) \in \Lambda_{[m]}$, thus $\mathscr{J} = \{(p, q, p, q) \in \Delta\}$. A similar argument is true for any cycle of even length.

(e) The graph is balanced bipartite and $(0, 0, 0, 0, \frac{1}{2}, \frac{1}{2})$ is a non-unstable equilibrium, since $\frac{\partial L}{\partial v_1} = -\frac{3}{5}$. By Remark 5.1, $\mathscr{J} = \{(0, 0, 0, 0, \frac{1}{2}, \frac{1}{2})\}$.

Future directions. The model introduced in [BBCL13] is more general than that defined in (1.1): fix $\alpha > 0$ and update the bins according to the rule

$$\mathbb{P}\left[i \text{ is chosen among } \{i, j\} \text{ at step } n\right] = \frac{B_i(n-1)^\alpha}{B_i(n-1)^\alpha + B_j(n-1)^\alpha}\,.$$

If $\alpha < 1$ then there is $w = w(G)$ such that $x(n)$ converges to w almost surely [BBCL13, Theorem 1.4]. For $\alpha = 1$ the present note and [BBCL13, CL13] establish convergence.

QUESTION 1: If $\alpha = 1$ and G is balanced bipartite, what is the distribution of the limit of $x(n)$?

In classical Pólya's urn the limit has a beta distribution, see [Pem07, Thm 2.1].

QUESTION 2: For $\alpha > 1$, does $x(n)$ converge almost surely?

QUESTION 3: For hypergraph-based Pólya's urns [BBCL13, §9.2], does $x(n)$ converge almost surely?

References

[BBCL13] M. Benaïm, I. Benjamini, J. Chen, and Y. Lima, *A generalized Pólya's urn with graph based interactions*, 2013. To appear in Random Structures Algorithms.

[Ben96] M. Benaïm, *A dynamical system approach to stochastic approximations*, SIAM J. Control Optim. **34** (1996), no. 2, 437–472.

[Ben99] ———, *Dynamics of stochastic approximation algorithms*, Séminaire de Probabilités, XXXIII, 1999, pp. 1–68.

[CL13] J. Chen and C. Lucas, *A generalized Pólya's urn with graph based interactions: convergence at linearity*, 2013. To appear in Electron. Commun. Probab.

[HJ13] R. A. Horn and C. R. Johnson, *Matrix analysis*, Second edition, Cambridge University Press, Cambridge, 2013.

[HPS77] M. W. Hirsch, C. C. Pugh, and M. Shub, *Invariant manifolds*, Lecture Notes in Mathematics, Vol. 583, Springer-Verlag, Berlin-New York, 1977.

[Pem07] R. Pemantle, *A survey of random processes with reinforcement*, Probab. Surv. **4** (2007), 1–79.

[Pem92] ———, *Vertex-reinforced random walk*, Probab. Theory Related Fields **92** (1992), no. 1, 117–136.

[SP00] B. Skyrms and R. Pemantle, *A dynamic model of social network formation*, Proceedings of the National Academy of Sciences of the United States of America **97** (2000), no. 16, 9340–9346.

Department of Mathematics, University of Maryland, College Park, MD 20742, USA.
E-mail address: yurilima@gmail.com